

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
11 October 2001 (11.10.2001)

PCT

(10) International Publication Number
WO 01/76194 A1

(51) International Patent Classification⁷: H04L 29/12,
12/24, 12/56

Alan [GB/GB]; 24 Foxglove Avenue, Needham Market, Ipswich, Suffolk IP6 8JJ (GB).

(21) International Application Number: PCT/GB01/00648

(22) International Filing Date: 16 February 2001 (16.02.2001)

(74) Agent: DUTTON, Erica, Lindley, Graham; BT Group Legal Services, Intellectual Property Department, 8th floor, Holborn Centre, 120 Holborn, London EC1N 2TE (GB).

(25) Filing Language: English

(81) Designated State (national): US.

(26) Publication Language: English

(30) Priority Data:
00302720.8 31 March 2000 (31.03.2000) EP

(84) Designated States (regional): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).

(71) Applicant (for all designated States except US): BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).

Published:


— with international search report

(72) Inventor; and

(75) Inventor/Applicant (for US only): BARRETT, Mark,

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: APPARATUS AND METHOD OF DETERMINING NETWORK ADDRESS USAGE AND ALLOCATION



(57) Abstract: The invention is concerned with a method of determining network address allocation on a communications network. The method comprises the steps of: b) obtaining one or more network addresses relating to one or more first nodes on the network for each first node: a1) obtaining a first list of network addresses of all nodes that are accessible via the first node for each node listed in the first list: a11) categorising the node; a12) if the categorised node corresponds to a network forwarding node, obtaining a second list, which second list comprises addresses that are accessible via the categorised node; a13) repeating steps a11-a12 until all network forwarding nodes that are accessible via the first node have been identified.

WO 01/76194 A1

APPARATUS AND METHOD OF DETERMINING NETWORK ADDRESS USAGE AND ALLOCATION

The present invention relates to a method of determining network usage and
5 allocation, suitable particularly, but not exclusively, for network management of Internet Protocol (IP) networks.

When administering both public and private address spaces, a network manager needs records of address space usage, spatial distribution of networks,
10 together with details relating to ownership and exact locations for individual addresses. This information is critical for the efficient running and security of the network. With IP networks each device needs to be allocated an address, and for IPv4, the allocation process is either static or dynamic. In both cases, the IP address space (either globally routable or private) needs to be managed so that, in particular
15 for globally routable address ranges, a company has proof that its currently allocated address space has been efficiently used. If it cannot prove this, it can be extremely difficult to convince the regional address registry that the network should be allocated more address space.

20 Currently companies use in-house methods to keep records of their address space and the input of data is typically a manual process. Maintaining current records, thus accounting for staff leaving etc., incurs major overheads, and is often not performed sufficiently frequently. Clearly, if the records are not correct, significant amounts of time can be wasted trying to locate a node with an IP address.
25 Furthermore, project-specific addresses are often not returned to the network manager at the end of a project, resulting in inefficient usage of network addresses and a network manager mistakenly claiming that the address space is exhausted.

Even if the manual records are considered to be fairly accurate by the
30 network manager, all that the records are likely to tell him is that an IP address was either last owned by X or part of a dynamic address pool. If the address was statically allocated then there may be some information on the location of the node but this information may or may not be valid. For dynamically allocated, or Dynamic

Host Configuration Protocol (DHCP), addresses there will be NO record of where this address is. In this case the network manager must interrogate his network (typically routers and switches) to determine the actual location of each IP address. In fault conditions this will lead to longer outages and require greater networking knowledge to locate rogue nodes. Incomplete or inaccurate records of IP usage can cause unnecessary levels of disruption to network users, particularly if the records concentrate on IP address allocation to hosts only. Consider the scenario of a host node located behind a switch on an Ethernet network, where the IP address of the switch has not been recorded by the network manager. If the host develops a fault, which affects other machines on the same network, then the network administrator is likely to disconnect all of the users on the network by disabling the corresponding router interface. If the switch had been recorded, however, then the network manager merely has to disconnect the port of the switch to which the host is connected.

15

Conventional network management tools are concerned with discovering the status of network nodes, together with the topology and the performance of the network, for example 3COM Transcend Central™, Cabletron Spectrum™ and Cisco CiscoView™. Many tools are available for collecting network traffic statistics, monitoring and logging errors and providing alarms and trends, for example NextPoint Networks NextPoint™ S3 2.5, Lucent Technologies VitalNet™ 7.0.

US Patent 5185860 describes a method for providing "Automatic discovery of Network Elements", where individual nodes are probed via network management protocols such as Simple Network Management Protocol (SNMP) and Internet Control Message Protocol (ICMP) echo requests commonly known as "ping host A". For a network with 1300 nodes, a considerable amount of time is spent in discovering and probing each node, and consequently the discovery process can only be run a few times each day. This, and other currently available discovery processes, is run once when discovering networks, and the discovered devices are then polled at periodic intervals in order to determine the status of the devices. In order to determine new networks, the whole process has to be run again.

According to a first aspect of the present invention, there is provided a method of determining network address allocation on a communications network comprising the steps of:

- a) obtaining one or more network addresses relating to one or more first nodes on
5 the network
- for each first node:
 - a1) obtaining a first list of network addresses of all nodes that are accessible
via the first node
 - for each node listed in the first list:
 - 10 a11) categorising the node;
 - a12) if the categorised node corresponds to a network forwarding
node, obtaining a second list, which second list comprises addresses that are
accessible via the categorised node;
 - a13) repeating steps a11 – a12 until all network forwarding nodes that
15 are accessible via the first node have been identified.

Preferably the first list of network addresses additionally comprises physical, or Ethernet, addresses corresponding to said network addresses. Each physical address is issued by one of said accessible nodes in response to a broadcast of a
20 network address from the first node. The process of broadcasting to, and receiving responses from, nodes that are accessible via the first node provides said corresponding physical addresses in the first list. The list therefore conveniently details physical and network addresses for each node, utilising the broadcast feature of the Address Resolution Protocol (ARP) to retrieve the physical address. This is
25 clearly beneficial over manual methods, which require updating each time a node is moved within a network and assigned a new network address.

Conveniently, information is obtained from nodes using a communications protocol, which is preferably the Simple Network Management Protocol (SNMP). This enables information to be retrieved from data maintained on the nodes themselves. In
30 addition, various performance-related information can be retrieved from the nodes via SNMP messages.

Advantageously the categorising performed at step (a11) includes identifying vendor specific information relating to the nodes. This can include transposing

network addresses of the categorised nodes into corresponding physical addresses, and comparing the first 3 bytes of the physical addresses with a plurality of predetermined physical addresses. This is a convenient means of identifying switch nodes, because Ethernet addresses give an indication of the vendor of the device,
5 which can be used to establish the type of node.

According to a second aspect of the present invention there is correspondingly provided apparatus for determining network address allocation on a communications network. The apparatus includes

- 10 (i) first means for obtaining one or more network addresses corresponding to one or more first nodes on a network;
- (ii) second means for obtaining, for each first node on the network, a first list of network addresses that are accessible via the first node;
- (iii) third means for categorising each node listed in the first list;
- 15 (iv) fourth means for obtaining, for nodes that have been categorised as network forwarding nodes by third means (iii), a second list of network addresses, which second list addresses correspond to nodes that are accessible via said network forwarding nodes; and
- (v) fifth means for inputting the second list to the third means (iii);
- 20 the apparatus being arranged such that the third, fourth and fifth means continue to operate until there are no more additions to the second list.

Each of the first, second, third, fourth and fifth means comprises a set of processable instructions to effect the functionality recited above.

25 Further aspects, features and advantages of the method of determining address allocation on a network will now be described, by way of example only as an embodiment of the present invention, and with reference to the accompanying drawings, in which:

Figure 1 is a schematic diagram of network devices whose network addresses are
30 discoverable by the method of the present invention;

Figure 2 is a schematic flow diagram describing a process of determining network addresses according to the present invention;

Figure 3 is a typical output of the method shown in Figure 2, showing IP addresses in use on a specified subnet; and

Figure 4 is a schematic block diagram showing in greater detail the processes present in a client and a server arrangement forming part of the embodiment of Figure 1.

5

In the following description, the terms "node", "device", "host" and "end host platform" are used. These are defined as follows:

"node": any equipment that is attached to a network, including routers, switches, repeaters, hubs, clients, servers; the terms "node" and "device" are used

10 interchangeably; and

"host": equipment for processing applications, which equipment could be either server or client, and may also include a firewall machine. The terms host and end host are used interchangeably.

15 *Overview*

It is generally recognised that IP address space needs to be managed. This is partly because companies are only allocated a finite number of addresses by a regional address registry, and partly because a well-managed address space enables more efficient trouble shooting. Currently each network device within network, e.g.

20 Local Area Network (LAN) or Intranet, which is engaged in sending and receiving communications, requires a network address. A network manager needs to monitor address usage so that if, for example, new devices are added to the network, he can identify which, and whether there are sufficient, network addresses available for allocation to each of the new devices.

25 Embodiments of the present invention are concerned with providing a method and apparatus for managing the address space in a network, such as the network shown in Figure 1. Figure 1 shows a generally conventional arrangement of a network 100, specifically an Ethernet type of network, having nodes comprising routers 101, switches 103 and hosts 105, interconnecting with a network 107 (only

30 one of each type of node has been labelled in Figure 1 for clarity). Nodes each have a physical address, or identifier, which identifies the node itself, and a network address identifying where it is in the network. In a conventional manner, a router 101 will make decisions on whether and where to forward packets that it receives on any of

its interfaces, usually based on the network address shown by the packet as a destination, possibly modifying the physical address or node identifier shown by the packet if required. Switches 103 interconnect multiple Ethernets, simultaneously transmitting multiple packets without modifying the packet, and hosts 105 are either
5 client or server machines (including database servers, web servers, proxy servers etc.) which run applications, some of which may transmit packets to, and receive packets from, other hosts on the network 100. Hosts 105 may also be firewall machines.

Specifically, embodiments of the present invention can be used to record and
10 interrelate network addresses, against address management data, by retrieving information from routing equipment such as switches and routers. Particularly conveniently, routing equipment such as routers and switches store address related data from packets that pass through them. Useful data for managing address space, which is not readily available elsewhere, can be obtained by accessing this stored
15 address related data. As a consequence of communicating with these devices only, the volume of network traffic generated is considerably less than for conventional systems, thereby enabling more frequent monitoring of the network. For example, for a network with 1300 nodes, comprising 10 routers and around 15 switches, a considerable amount of time and network traffic is saved when information is
20 gathered from these routing devices only, instead of all 1300 nodes.

Although it is known that network address data is available on router devices, in the form of Address Resolution Protocol (ARP) tables, advantageous embodiments of the invention utilise a novel method for identifying and locating switches and understanding cascaded switch connections, as described in detail
25 below (The term "cascaded switch" indicates that that two or more switches are connected, either in a series or a parallel arrangement, to an interface of a router (or other network device)).

Embodiments of the invention also include means for discovering Virtual Local Area Networks (VLANs) that are configured on switches (VLAN is a local area
30 network where nodes are mapped according to a criterion other than geographic location (for example, by department, type of user, or primary application)). This feature is novel over known topology/discovery-related methods, which discover network topology using data collected from routers alone. Such existing methods are

limited in the scope of their discovery, because VLAN topology on switches cannot be identified from router data. Thus known methods, which base their discovery process on data retrieved from routers, are unable to identify VLAN configurations. This is a problem because the configuration of VLANs affects the bridging
5 (forwarding behaviour) of the switches and can therefore affect the reachability of a node. Thus a particularly advantageous feature of the invention is the ability to discover hosts that are only reachable with knowledge of VLAN configuration.

Yet further advantages of embodiments of the invention lie in the way in which data collected from nodes is collected and used: in particular data is collated
10 by subnet (nodes within a given range of IP addresses; router (nodes directly connected to each interface of the router); switch (nodes directly connected to that switch – connection identified by slot and port number); VLAN (nodes connected to a specified logical segment); and Single IP or MAC address, or Domain Name System (DNS) name (the DNS corresponding to the IP address in question). This information
15 is then used to identify unused network addresses that have been allocated by the regional address registry, but have not been assigned to a network device, and this allows network managers to manage address space more efficiently.

Particularly advantageously, the information gathered by embodiments of the invention can be used to improve the process of dynamic address assignment.
20 Dynamic address assignment typically uses the Dynamic Host Configuration Protocol (DHCP) to distribute a new IP address when a device connects to a different place in the network. Typically, there is NO record of where DHCP addresses are, and a network manager has to manually interrogate his network (typically routers and switches) to determine the actual location of each IP address, which can be
25 extremely time consuming. As outlined above, and as described in detail below, embodiments of the invention continually monitor allocation and use of IP address locations, thereby enabling the network manager to easily identify the location of IP addresses.

30 An embodiment of the invention, generally referred to as address determining apparatus 109 for performing the method of determining network addresses on a network 100, may be stored on the hard disc drive of a host machine, shown as host 106 in Figure 1, for processing thereon. The address determining apparatus 109 may

be located on any host machine (implementation details given later), providing the apparatus can access each of the routers 101 on the network. The method described below retrieves information relating to network nodes by issuing SNMP messages to a Management Information Base (MIB) that is maintained on the node. SNMP, or
5 Simple Network Management Protocol, is part of the known TCP/IP network software, and MIB, or Management Information Base, is a standard specifying the data items that a host, router or switch must keep, together with the operations allowed on each. SNMP is the protocol that enables information to be extracted from a MIB, and is known to those skilled in the art. For further details see Request for
10 Comments (RFC) 2037/2737, Entity MIB, McCloghnie *et al* 1996/1999, published by the Internet Engineering Task Force (IETF) (available from <http://www.ietf.org>), or Understanding SNMP MIBs by David Perkins, Evan McGinnis. Prentice Hall, 1st edition (December 3, 1996);

15 *Method for determining network address usage*

Figure 2 shows a flow diagram of a method according to an embodiment of the present invention, which, as described above, can be run from a host 106 of Figure 1, with the precondition that all routers attached to networks to be managed are accessible to the address determining apparatus 109, and that the network
20 addresses of these routers are predetermined.

- S 2.1 Retrieve a first router network address from a predefined list of router addresses – e.g. from a file;
- S 2.2 Send an SNMP message to the router, requesting access to the router's MIB (Management Information Base), in particular to the corresponding router
25 ARP (Address Resolution Protocol) table.
- S 2.3 If available, retrieve any CDP (Cisco Discovery Protocol™) information stored on the router. CDP is a media- and protocol-independent device-discovery protocol that runs on all Cisco-manufactured equipment including routers, access servers, bridges, and switches. Using CDP, a device can advertise its existence to other
30 devices and receive information about other devices on the same LAN or on the remote side of a WAN (Wide Area Network). These devices broadcast information about themselves to neighbouring devices via packets, which are stored on these devices for discovery via SNMP, or Telnet. There will therefore only be CDP

packets if the router is a Cisco™ type of router, and if any neighbouring nodes are Cisco™ devices.

- S 2.4 If there are CDP packets then this information relates to directly connected Cisco™ devices; store the CDP information in a Cisco™ linked list (or log file), and
5 send further SNMP messages to each of the devices to retrieve various operational parameters, forwarding tables and ARP tables if available (described in greater detail below);
- S 2.5 Filter, or remove, the Cisco™ Ethernet addresses from the ARP table retrieved from the router, so as to generate a modified router ARP table (modified
10 table thus contains IP addresses of non-Cisco™ nodes). Note that if the router itself is not a Cisco™ router, steps S 2.3 and S 2.4 will be redundant, as the router will be impassive to any CDP broadcasts. Thus the method will skip steps S 2.3 and S 2.4, and the modified ARP table = retrieved ARP table;
- S 2.6 For each of the entries in the modified ARP table, inspect the first three
15 Bytes of the Ethernet address in order to determine whether the address matches a known device (router and/or switch) vendor allocation. If it does then issue SNMP messages to discover various operative parameters relating thereto, retrieving a forwarding table and an ARP table (described below), if available, and storing these parameters in a non-Cisco™ log file;
- 20 ♦ If the device corresponding to the first network device either has an address of a known vendor, or doesn't respond to SNMP then mark the device as an end host platform and save in end host log file;
- S 2.7 For each network addresses listed in the Cisco™ log file, repeat steps S 2.3 to S 2.6;
- 25 • S 2.8 For each network address listed in the non-Cisco™ log file, repeat step S 2.6 (as these network addresses will be non-Cisco™ devices, these devices are not operable to receive CDP packets);
- S 2.9 Repeat steps S 2.7 and S 2.8 until all network addresses that are connected, directly or indirectly, to the first router have been determined;
- 30 • S 2.10 Retrieve a second router network address from the predefined list of router addresses and repeat steps S 2.2 – S 2.9;
- S 2. 11 Repeat S 2.10 until all router network addresses have been traced (not shown).

Inspection of CDP packets and SNMP messaging:

If information is gathered from a Cisco™ device that supports CDP, the following information is available in response to a SNMP message, or a Telnet to the
 5 router, eg.:

prompt# Telnet 10.10.10.1

router# enable

router# show cdp neighbors

OUTPUT:

Device ID	Local IF	Hold_t	Capability	Platform	Port ID
066510624(Topcat)	ATM0.2	158	T B S	WS-C5505	5/1
001936193(Brains)	ATM0.10	123	T S	WS-C5000	3/1
001936193(Brains)	ATM0.3	123	T S	WS-C5000	3/1
003272985(pumyra)	ATM0.2	161	T S	WS-C5000	5/1
014580358(FTB)	ATM0.3	167	T B S	WS-C5000	4/1
014580358(FTB)	ATM0.4	167	T B S	WS-C5000	4/1

10 Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge

S - Switch, H - Host, I - IGMP, r - Repeater

TABLE 1

"Capability" defines what this device is capable of achieving, e.g. bridge
 15 (switch at layer 2), router (layer 3) etc, and the device type can be determined from this information. The MIB information varies with device type, such that once the device type has been established, additional information can be collected via appropriate SNMP messages. Switches that bridge Local Area Network (LAN) segments are obliged to maintain a MIB standard specified in the IEEE 802.1D-1990,
 20 defined in RFC 1493, Definitions of Managed Objects for Bridges, Decker *et al*/ July 1993 (available from <http://www.ietf.org>). This is essentially a database comprising information relating to the switch, and the Bridge-MIB defines a minimum standard that switch manufacturers must support. Cisco™ switches maintain a proprietary MIB, known as Cisco-CDP-MIB, which includes additional features that are unique to
 25 Cisco™ devices of this type. Further information is available from <http://www.cisco.com/public/mibs>.

Considering a scenario of a router being connected to a switch, SNMP uses the device identity (ID) (see table 1 above) to determine:

- a) device IP address;
- b) local port (interface on local device (router) used by neighboring node (switch);
- 5 c) remote port (interface used by remote device (e.g. switch port) to get to the local device (router)) and
- d) Platform (specifies type of device) - although this can be derived from the capability code from the MIB.

For non-Cisco™ devices, and as described at step S 2.6, the corresponding
10 Ethernet addresses give an indication of the vendor of the device: Ethernet hardware manufacturers purchase blocks of Ethernet addresses and assign them in sequence as they manufacture Ethernet interface hardware. Assuming the address determining apparatus 109 has access to a list of Ethernet addresses (typically the first 3 Bytes thereof), the Ethernet addresses can be compared with such a predetermined list,
15 thus providing a first estimate as to whether or not the corresponding device is likely to be a switch. If the Ethernet address corresponds to a switch vendor, SNMP is then used to probe the device further to access, for example, the Make and Model of the switch. This information is stored in the Bridge-MIB described above, specifically in a MIB parameter known as System Object Identity. Other parameters are similarly
20 retrievable from the Bridge-MIB, including those listed above for the Cisco™ devices.

Discovery of Indirectly Connected Devices

Steps S 2.7 to S 2.9 describe the discovery of devices that are indirectly connected to the router. In the present embodiment, the network addresses of all of
25 the routers (starting points) are explicitly defined as described in step S 2.1 - e.g. in a file. Therefore the only unidentified network devices that are discoverable by multi-layer probing are switches and host machines. ("Unidentified" is used here to mean simply a connected network device whose network address was not amongst those explicitly defined (thus not a router) and has not yet been discovered.) Host
30 machines will either have Ethernet addresses that correspond to non-switch vendors, or will not respond to SNMP messages; thus the remaining devices to be discovered are expected to be switches only.

Switch devices are determined by matching the vendor part of the Ethernet

address against entries in the switch log file. For Devices identified as switches and not previously discovered via CDP (as previously described), the address determining apparatus 109 can access a forwarding table that is internally maintained by the switch. Such a forwarding table lists Ethernet address and port number for packets
5 seen by the switch, specifying the switch ports used to relay the packets. Thus the passage of packets through a switch interface will create an entry in the switch's forwarding table, enabling the identification of devices connected to the switch.

Packets that pass through switches are typically destined for a host machine only. Consider the scenario of a first switch being directly connected to a router, and
10 having a second switch connected via one of its ports. As a switch is not an end host, packets passing through the first switch are unlikely to have a destination address corresponding to a switch. Thus neither the ARP table of the router nor the forwarding table of the first switch would include the Ethernet addresses of the second switch. However, providing that a network address is manually configured
15 when the second switch is installed, the switch will use ARP to determine the MAC address of its default gateway. This process results in a valid router ARP table entry plus a forwarding table entry within the first switch that corresponds to the second switch. Embodiments of the invention are configured to capture the ARP table entry and forwarding table entry shortly after the switch has broadcast to its default
20 gateway, so that the Ethernet address of the switch can be recorded by the address determining apparatus 109.

Information gathered

As stated above, the method makes use of the fact that routers and
25 switches maintain records of packets (thus destination addresses of nodes) that pass through them. Hence, if a node is in use (receiving data and/or sending data), its IP address (and Ethernet address) will be stored in the router or switch (for a predetermined length of time). The above method accesses information gathered from CDP packets and/or ARP or forwarding tables, and these sources provide a
30 basis for gathering further information relating to the nodes (discovered IP addresses). This further information is gathered using SNMP messages, and includes the following:

FOR EACH IP ADDRESS

Information	Source
<p>Media Access Control (MAC) address:</p> <p>For end devices this field shows the Ethernet address. Network and Broadcast addresses are identified, together with addresses owned by a router.</p>	<p>ARP tables in router and bridging tables in switches</p>
<p>First hop network device:</p> <p>If the first hop device is a switch, the slot/port numbers (the interface this IP address is connected to – see below).</p> <p>If the first hope device is a router, the Interface that the IP address is connected to.</p>	<p>SNMP on router: When collecting ARP table information from a router, each entry is linked with its first hop router. If the device is then discovered to be behind a switch, then the first hop field is over written with the IP address of the switch</p>
<p>Switch slot/port:</p> <p>Switches normally have multiple slots into which interface modules are fitted. Each interface module will have 1 to N interfaces (ports). Slot therefore shows which module the address is connected to and the port refers to the physical position. The slot/port combination will therefore provide a unique reference point</p>	<p>SNMP on bridge MIB: search for the MAC address that was retrieved from the router ARP table on the switch. If the MAC address is active then it will appear in the dynamic bridge MIB table on the switch.</p>
<p>VLAN member (if appropriate):</p> <p>Defines which logical segment the IP address belongs to, if there are logical segments that are isolated from each other.</p>	<p>SNMP on switch forwarding table (switches maintain a separate bridging (layer 2 forwarding table) per VLAN).</p>
<p>Date last seen:</p> <p>Time that a packet destined for this IP address was last seen.</p>	<p>When run, the apparatus 109 obtains a time stamp from the operating system. All modified</p>

	IP records place this time stamp in the date last seen field.
--	---

TABLE 2

As is known in the art, the MIB maintains statistics relating to a range of
 5 network parameters, and the above table provides an example only (non-exhaustive list) of data that may be collected according to the present invention.

Collating Information gathered

Each IP address record is stored in a file in the format shown below (each file
 10 goes from address 0 to 255). The advantage of using integers is that, for example, an IP address is stored in 4 bytes rather than 16 bytes if stored in ASCII. This reduces the file sizes (saving disk space) and the structure provides fixed fields to allow fast, efficient and reliable parsing of data. It is understood, however, that it is inessential to the invention to store data in this format.

15

```

struct {
    unsigned long int ip;      /*IP address of node*/
    char mac[18];             /*Physical address corresponding to IP address*/
    unsigned long int upstream; /*IP address of first hop router or first hop switch*/
    20 unsigned long port;      /*Interface number on switch or router to which the node
                               is connected, retrieved via SNMP interface tables for
                               routers or for switches (Cisco) */

    int date;                 /*Time stamp from OS*/
    unsigned long int hub;    /* Flag indicating that node is connected, or not, to a
    25 hub; flag takes different values depending on whether
                               node relates to a end-host address that is not connected
                               to a hub; an end-host address that is connected to a hub;
                               a network address; a broadcast address; a reserved
                               address etc */

    30 int vlan;               /*logical segment to which this node is connected*/
  
```

15

```

char info[51];          /* text to assist the administration process (typically
                        manually added) */
} file_data[257];

```

- 5 With data being available in this format, information (details as per variables within file_data) can be displayed as a function of:
- ❖ subnet (nodes within a given range of IP addresses, typically for a C class subnet (each IP address is a pair (network, host on network); for a class C address, 21 bits are allocated to the network, and 8 bits to the hostid – e.g. a class C subnet
 - 10 may be 132.146.107.0 – 132.146.107.255));
 - ❖ router (nodes directly connected to each interface of the router);
 - ❖ switch (nodes directly connected to that switch – connection identified by slot and port number);
 - ❖ VLAN (nodes connected to a specified logical segment)
 - 15 ❖ Single IP or MAC address, or Domain Name System (DNS) name (the DNS corresponding to the IP address in question)

Figure 3 shows a typical output for nodes located on subnet 132.146.107.0.

Implementation

- 20 As described with reference to Figure 1, address determining apparatus 109 to effect the method of the above embodiment may be loaded on a client terminal 106. The apparatus 109 can be run by a user, for example a network manager or network administrator, to assess current address usage. The user enters data via a browser, which provides a form for the user to specify a request in a known manner.
- 25 Referring to Figure 4, stored within the client terminal 106 (e.g. on the hard disk drive thereof) is an operating control program 410 comprising an operating system 412 (such as Windows (TM)), a browser 414 (such as Netscape (TM)) and application 411a, designed to operate within the browser 414. The function of the operating system 412 is conventional and will not be described further. The function
- 30 of the browser 414 is to interact, in known fashion, with hypertext information 411a received from a server 420 via a LAN (the server 420 may be one of the hosts 105 shown in Figure 1). In this embodiment the hypertext information may be an HTML form, which is displayed to the user. The user then enters various parameters and/or

requests, and posts the form, in a known manner, to a co-operating program 411b; thus form 411a and co-operating program 411b comprise the address determining apparatus 109. This form essentially captures any parameters entered by a user and transfers them to the co-operating program 411b stored on the server 420. For
5 further information see "Client/Server Programming with Java and Corba", 2nd Edition, R. Ofrali and D. Harkey, pp. 239 - 242. Typical parameters that are entered by the input HTML form include a list of routers for steps S2.1, S2.10 and S2.11, together with a field from which to display the data – e.g. subnet, router, switch etc.

The co-operating program 411b, having received the completed form 411a,
10 writes the input data to a configuration file 422 stored on the server 420. This is convenient if the co-operating program 411b is to be run several times during the day for the same input parameters; in this situation data collected is stored in an output file 424, for subsequent display as required (i.e. it is not immediately posted to an output HTML form). The co-operating program 411b acts on data in the configuration
15 file 422 according to the method presented in Figure 2. In order to send and receive information to and from routers as described above, the co-operating program 411b connects to each of the routers 101 shown in Figure 1 in a known manner via sockets. Once the co-operating program 411b has carried out this method, the collated information is inserted into a reply HTML document and displayed to the user
20 on the browser 414 (e.g. Figure 3).

It is understood that the use of HTML forms in this manner is inessential to the invention: an application to effect the method of the embodiment could be stored on the server as an applet, downloaded into the browser 414, and run within the browser 414 in a known manner. Alternatively the method could be embodied in a
25 Windows™ - based application loaded on the client terminal 106.

The co-operating program 411b is written in the C-programming language and pre-configured with directory location of the configuration files and the directory location for the data records. When the co-operating program 411b is loaded on a
30 unix server, it can be automatically run by making a Cron entry (Cron is a job scheduler unix utility). This allows the user to run the capture program as frequently as is deemed necessary to ensure that the records are up to date.

Example Cron entry:

O 6,8,10,12,14,16,18 * * * /home/barretma/ip-rec-config/ip-rec-test

In the example above the co-operating program 411b "ip-rec-test" is started in the first minute of the hours 6,8,10,12,14,16,18 of every day.

5

Other Advantages

DHCP (Dynamic Host Configuration Protocol) is a protocol that allows centralised automation of the assignment of Internet Protocol (IP) addresses in a network. When an organisation sets up its computer users with a connection to the Internet, an IP address must be assigned to each machine. DHCP allows distribution of new IP addresses when a computer is plugged into a different place in the network. For these DHCP addresses there is typically NO record of where these addresses are. Thus a network manager must manually interrogate his network (typically routers and switches) to determine the actual location of each IP address, and this can be extremely time consuming. By carrying out the above embodiment, IP address locations (in terms of first hop router, for example) and usage can be continually monitored.

This present invention is thus a proactive network address space management method, and provides a convenient and efficient alternative to the reactive methods that are currently employed by most network administration programs.

Modifications

The list of routers currently entered via the HTML form 411a or via configuration file 422 may alternatively be stored on a computer that is remote from the client terminal 106 and the server 420. In this case, the method would include a further method step of retrieving the file from the remote location. As an alternative to explicitly listing routers to be probed, they could be automatically discovered according to the method described above for discovering switches. In this case the user would have to specify a minimum of one router as a starting point.

30

It is mandatory for all vendors of switches and routers to maintain a MIB, but it is not mandatory for them to support SNMP. In the event of the network including devices that do not support SNMP, the above method could pass messages to the

MIBs of the routers and switches using Telnet. The co-operating program 411b would thus embed "CHAT" like scripts, which issue a series of Telnet messages to the devices. However, commands often vary between vendors, and in order to effect communication between the program and devices, the vendor and device Operating
5 System version is needed. Once this information is known, the scripts can be written using appropriate syntax, and messages parsed from program to devices. The format of replies received in response to the queries also varies between different vendors, so the program will have to include means for understanding a variety of reply formats.

10

The method described above probes each and every device for which an IP address is registered in the ARP table of the router, and retrieves substantially the same information for each type of device. However, the method could be modified to initially retrieve the object ID only (via SNMP), and, dependent on this object ID,
15 process a predetermined series of SNMP messages (via a pre-configured rule-set). This therefore provides a tailored extraction of information depending on the device.

The above embodiment describes retrieving information from routers and switches only. However, DHCP servers, which maintain a lease file (which is much
20 like a router ARP table in that Ethernet address to IP mappings are known), provide another source of address usage information. Thus the lease file could be interrogated.

The present method describes saving data into log files; however, for storage
25 of IP records on a very large scale where performance and resilience is essential, retrieved data should preferably be saved on a commercial database like Oracle.

Other information, in particular performance statistics relating to the devices, is retrievable from a device's MIB. As a simple extension to the above method, data
30 such as device timeouts, unavailability and usage thresholds, could be collected via SNMP. This could be used to implement performance traps between the output of the co-operating program 411b and a Network Management station typically located in a network operations, to enable proactive monitoring of network device behaviour.

The above embodiment describes operation of the invention for an Ethernet LAN, but the invention is operable on any IP network, due to the layered architecture of networks. Thus the invention may also be applied to Asynchronous Transfer Mode
5 (ATM) networks between routers, ATM LANE networks between routers and switches and Token Ring network, among others.

As will be understood by those skilled in the art, the invention described above may be embodied in one or more computer programs. These programs can be
10 contained on various transmission and/or storage mediums such as a floppy disc, CD-ROM, or magnetic tape so that the programs can be loaded onto one or more general purpose computers or could be downloaded over a computer network using a suitable transmission medium.

CLAIMS

1. A method of determining network address allocation on a communications
5 network, the method comprising the steps of:
 - a) obtaining one or more network addresses relating to one or more first nodes on the network
for each first node:
 - a1) obtaining a first list of network addresses of all nodes that are accessible
10 via the first node
for each accessible node listed in the first list:
 - a11) categorising the accessible node;
 - a12) if the categorised node corresponds to a network forwarding node, obtaining a second list, which second list comprises addresses that are
15 accessible via the categorised node;
 - a13) repeating steps a11 – a12 until all network forwarding nodes that are accessible via the first node have been identified.
2. A method according to claim 1, further including the step of storing the first and
20 second lists of addresses, together with their categorised types of nodes.
3. A method according to claim 1 or claim 2, in which the first list of network addresses additionally comprises physical addresses corresponding to said network addresses, each of which physical addresses is issued by one of said
25 accessible nodes in response to a broadcast of a network address from the first node, which process of broadcasting to, and receiving responses from, nodes that are accessible via the first node provides said corresponding physical addresses in the first list.
- 30 4. A method according to anyone of the preceding claims, in which the categorising includes identifying vendor specific information relating to the nodes.

5. A method according to claim 4, in which said identification of vendor specific information includes the steps of:
- (i) transposing network address into a physical address; and
 - (ii) inspecting the first 3 bytes of the transposed physical address and comparing
- 5 the same with a plurality of predetermined physical addresses.
6. A method according to any one of the preceding claims, including the steps of
- (i) accessing categorised network forwarding nodes so as to obtain packet statistics information relating to end-host nodes connected thereto; and
 - 10 (ii) adding the packet statistics information to network address corresponding to the end-host nodes in the respective first or second lists.
7. A method according to any one of the preceding claims, in which information is obtained from nodes using a communications protocol.
- 15
8. A method according to claim 7, in which the communications protocol is the Simple Network Management Protocol.
9. Apparatus for determining network address allocation on a communications
- 20 network including
- (i) first means for obtaining one or more network addresses corresponding to one or more first nodes on a network;
 - (ii) second means for obtaining, for each first node on the network, a first list of network addresses that are accessible via the first node;
 - 25 (iii) third means for categorising each accessible node on the list and thereby identifying network forwarding nodes among said accessible nodes;
 - (iv) fourth means for obtaining, for nodes that have been identified as network forwarding nodes by third means (iii), a second list of network addresses, which second list comprises nodes that are accessible via said network forwarding nodes;
 - 30 and
 - (v) fifth means for inputting the second list to the third means (iii);
- the apparatus being arranged such that the third, fourth and fifth means continue to operate until there are no more additions to the second list.

10.A method of listing network address allocation on a communications network, which network address allocation has been determined according to claim 1, the method including the steps of

- 5 (i) specifying a first network identifier ;
- (ii) identifying, from the first and second lists, any network addresses that are directly connected to the first network identifier;
- (iii) retrieving any information corresponding to the identified addresses that is included in the first and second lists for displaying to a user.

10

11.A method according to claim 10, wherein the network identifier includes any one of

- (i) a router network address;
- (ii) a subnet network address;
- 15 (iii) a switch network address;
- (iv) a VLAN number;
- (v) an end host network address.

12.A computer program comprising a set of instructions to cause a computer to
20 perform the method according to any one of claims 1-8.

13. Address management apparatus for use in managing network address allocation in a communications network, the network comprising nodes connected by communications links,

25 at least a first of which nodes being provided with a routing information data store for use in routing data traffic over the network to at least one destination node, said data traffic including address-related information for destination nodes in the network,

said first node further comprising means to access and copy said address-related information to an address-related information data store,
30 wherein the address management apparatus comprises means to access the address-related information data store and to retrieve address-related information therefrom for use in managing network address allocation in the communications network.


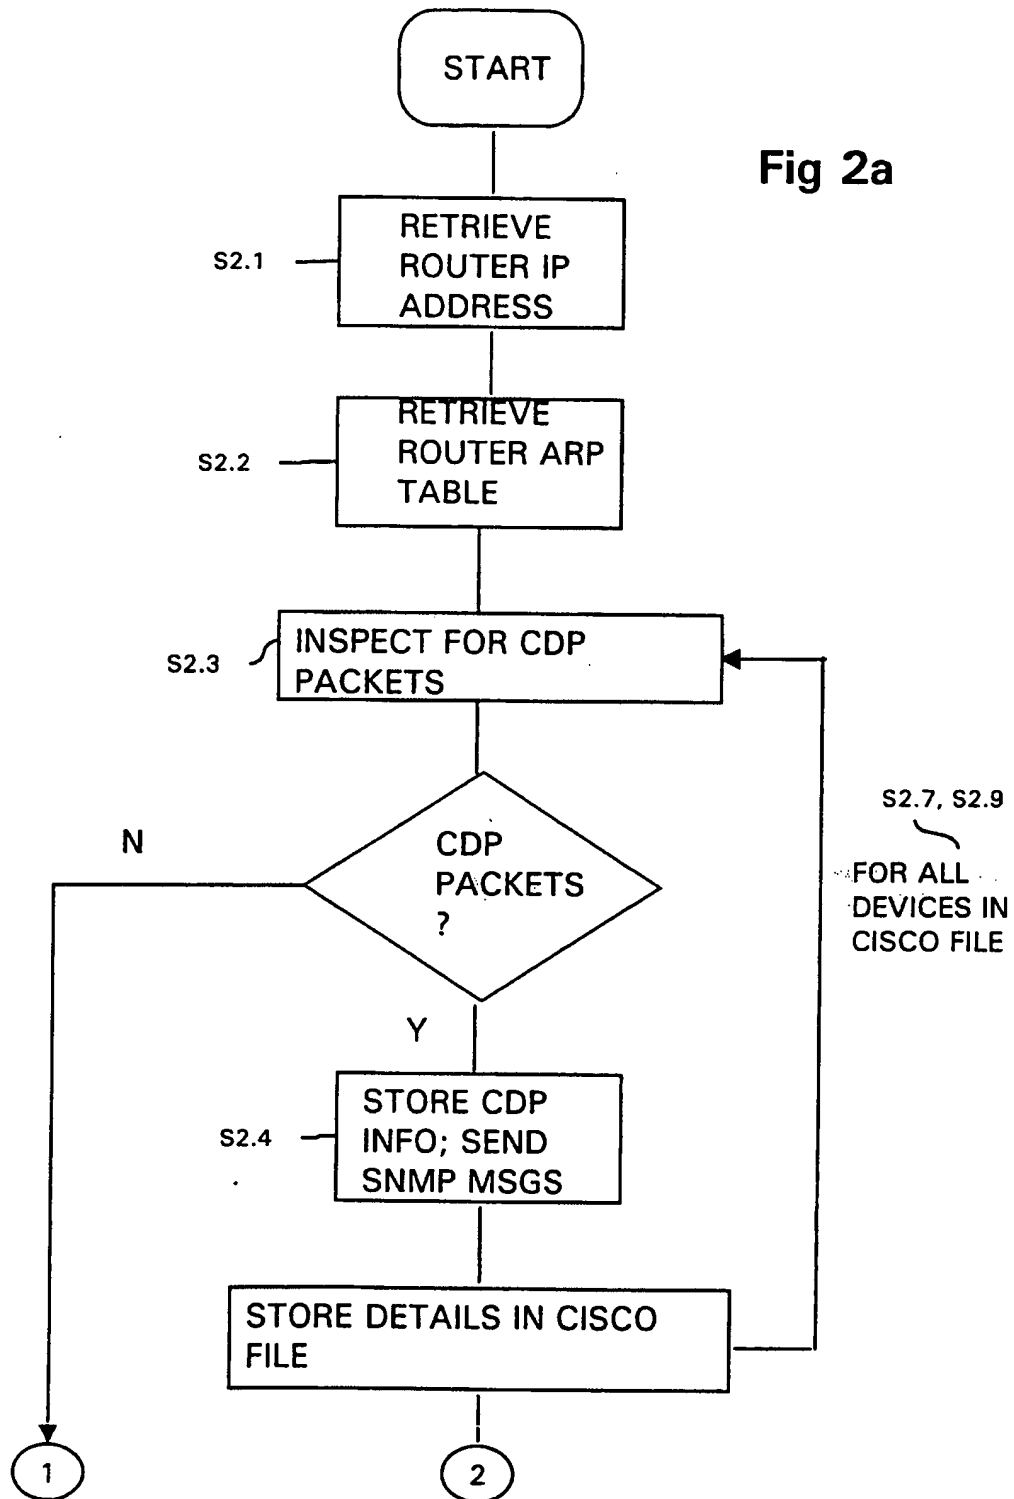


Fig 1

2/4

Fig 2a



3/4




Fig 2b

4/4

IP Address	MAC Address	Location	VLAN ID	HUB	Date (last seen)	Information
132.146.107.0	-	churchill			02.03.00:14:00	
132.146.107.1	Local	churchill			02.03.00:14:00	Churchill: B54 comms rm
132.146.107.2	00-40-0B-58-5F-FF	mufsa	3		02.03.00:14:00	Mufsa: B54 frame room
132.146.107.3	00-10-1F-2B-B4-FF	churchill			02.03.00:14:00	Garfield
132.146.107.4	00-40-0B018-63-FF	mufsa	3		02.03.00:14:00	Bagpuss

Fig 3




Fig 4

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 01/00648

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 H04L29/12 H04L12/24 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, IBM-TDB, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y A	EP 0 809 383 A (SUN MICROSYSTEMS INC) 26 November 1997 (1997-11-26) abstract page 2, line 33 -page 3, line 3 page 3, line 49 -page 4, line 4 page 5, line 11 -page 5, line 18 page 4, line 44-56 page 6, line 20-45 --- Y A EP 0 511 851 A (HEWLETT PACKARD CO) 4 November 1992 (1992-11-04) abstract page 3, line 12-21 page 4, line 55 -page 5, line 11 page 6, line 27 -page 7, line 3 claim 1 --- -/-	1-3,7,9, 12,13 6 4,5,8, 10,11 6 1-5,7-13

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

& document member of the same patent family

Date of the actual completion of the international search

30 March 2001

Date of mailing of the international search report

06/04/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040. Tx. 31 651 epo nl.
Fax (+31-70) 340-3016

Authorized officer

Cichra, M

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 01/00648

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 796 736 A (SUZUKI HIROSHI)	1, 3, 9, 13
A	18 August 1998 (1998-08-18) abstract	2, 4-8, 10-12
	figures 11, 14, 16 column 4, line 32 -column 5, line 45 ---	
A	US 5 588 119 A (HARE DUNCAN ET AL) 24 December 1996 (1996-12-24) claims 1-4 column 6, line 23-60 -----	1-13

INTERNATIONAL SEARCH REPORT

Information on patent family members

I. International Application No

PCT/GB 01/00648

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0809383 A	26-11-1997	US 5835720 A JP 10056451 A	10-11-1998 24-02-1998
EP 0511851 A	04-11-1992	US 5297138 A CA 2061687 A DE 69225637 D DE 69225637 T JP 5183549 A	22-03-1994 31-10-1992 02-07-1998 24-09-1998 23-07-1993
US 5796736 A	18-08-1998	JP 2871469 B JP 8032597 A CA 2154099 A DE 19526001 A	17-03-1999 02-02-1996 20-01-1996 01-02-1996
US 5588119 A	24-12-1996	NONE	